Power Systems

IBM

Coherent Accelerator Processor Interface (CAPI)

Overview Package

CAPI Developer Kit

What's in a name?

Coherent

Accelerator

Processor

Interface

FPGA as an Accelerator

Accelerator

- FPGA: Field Programmable Gate Array
  - It's a re-programmable chip
  - It can run fast (cycle times of 250 – 500 Mhz or more)
  - It has Industry Standard Interfaces like PCI-E Gen3
  - The Major FPGA Suppliers, Altera and Xilinx, are OpenPOWER Foundation members

ALTERA. XILINX.

Source code for FPGAs has traditionally been written in RTL* (VHDL** or Verilog). Now, we also have OpenCL, a more programmer friendly language.

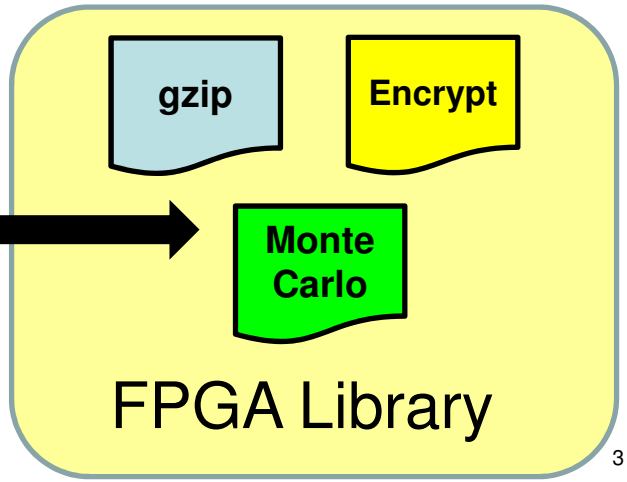

*RTL = Register Transfer Level
VHDL = VHSIC* Hardware Description Language
***VHSIC = Very High Speed Integrated Circuit

When to Use FPGAs

- Transistor Efficiency & Extreme Parallelism
  - Bit-level operations
  - Variable-precision floating point
- Power-Performance Advantage
  - >2x compared to Multicore (MIC) or GPGPU
  - Unused LUTs are powered off
- Technology Scaling better than CPU/GPU
  - FPGAs are not frequency or power limited yet
  - 3D has great potential
- Dynamic reconfiguration
  - Flexibility for application tuning at run-time vs. compile-time
- Additional advantages when FPGAs are network connected ...
  - allows network as well as compute specialization

Why is an Accelerator Faster?

Question: The POWER8 Processor runs at ~3Ghz while our FPGA runs at 250Mhz. So why would an accelerator be better?

Answer: The FPGA is better for certain algorithms, such as those that are numerical intensive or have parallelism.
The POWER8 processor has a finite set of instructions to implement the algorithm in SW.
The FPGA is customized logic built for specific processing of an algorithm.

5

Why is an Accelerator Faster?

Accelerator

PCIE FPGA

Example 1: Numerical Intensive Algorithm

$$\int P(x)dx = a_0 + \sum_{n=1}^{k} \left(a_n \cos \frac{n\pi v}{L} + a_n \sin \frac{n\pi w}{L} \right)$$

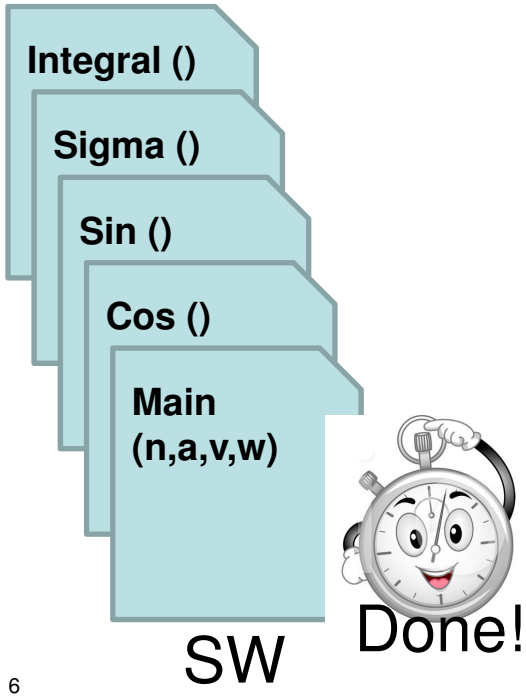

Variables

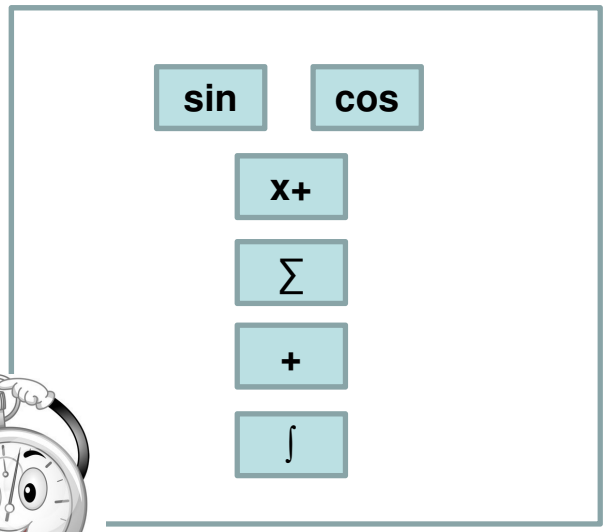

6

Why is an Accelerator Faster?

Accelerator

PCIE FPGA

Example 2: Parallelism

Monte Carlo Risk Analysis to determine probability of financial success:
Given current finances, run 100 scenarios

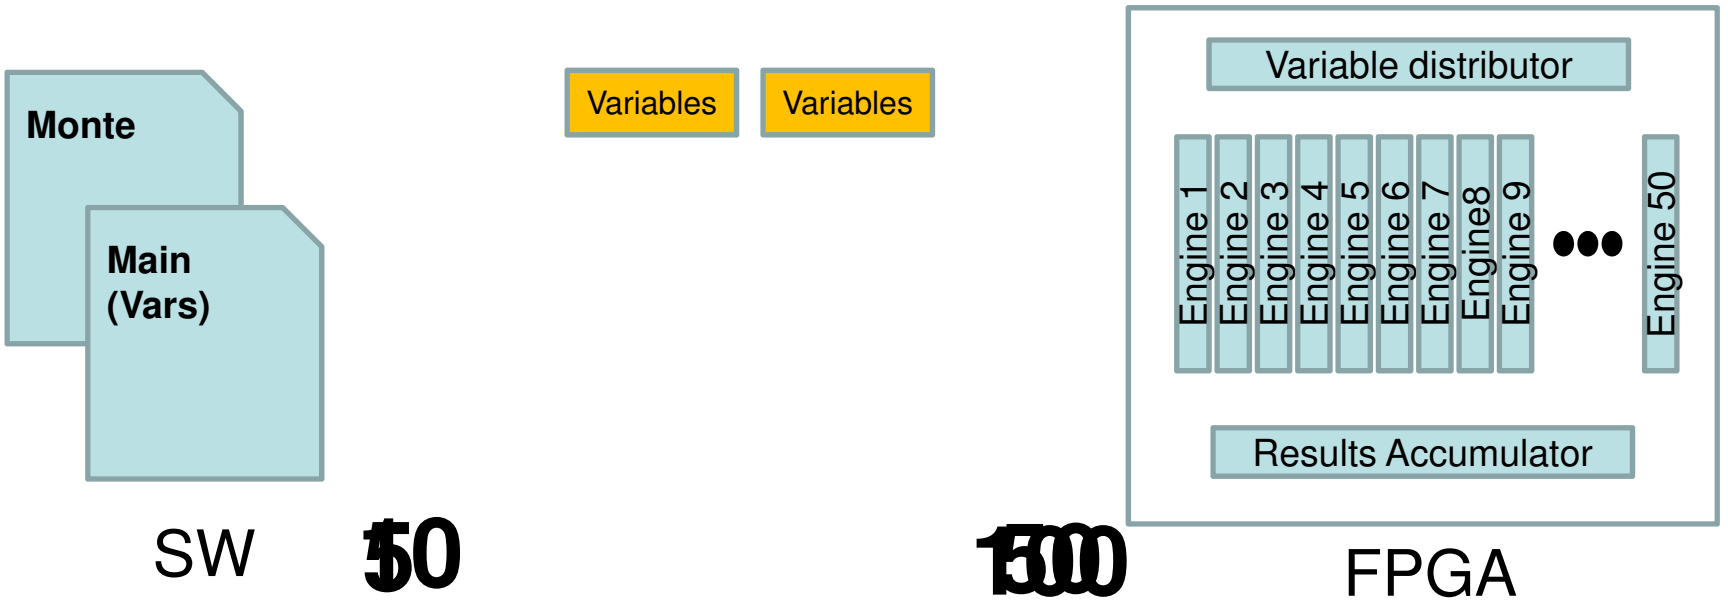

7

So what is new?

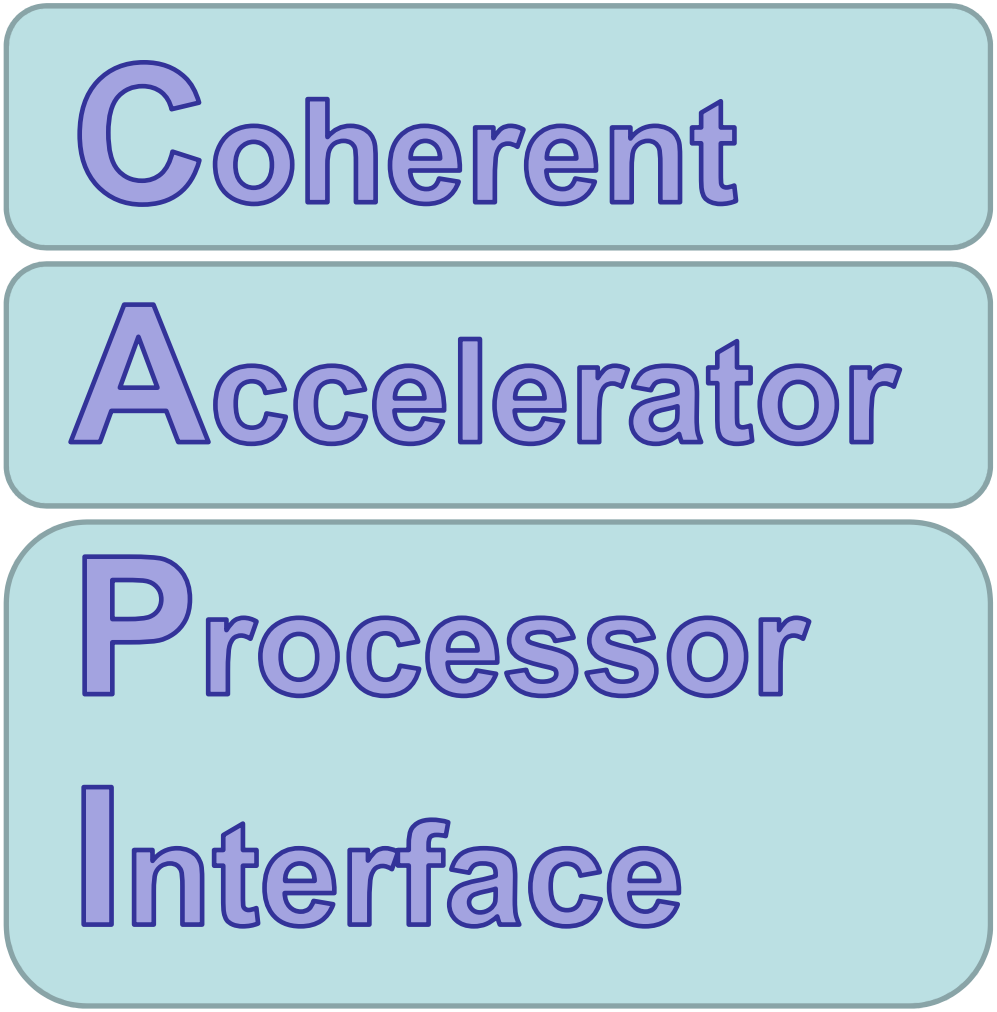

Accelerators on FPGAs have been around for a long time….
So what is new?

Coherency makes the accelerator a peer to the POWER8 cores

What was done before CAPI?

Prior to CAPI, an application called a device driver to utilize an FPGA Accelerator.
The device driver performed a memory mapping operation.

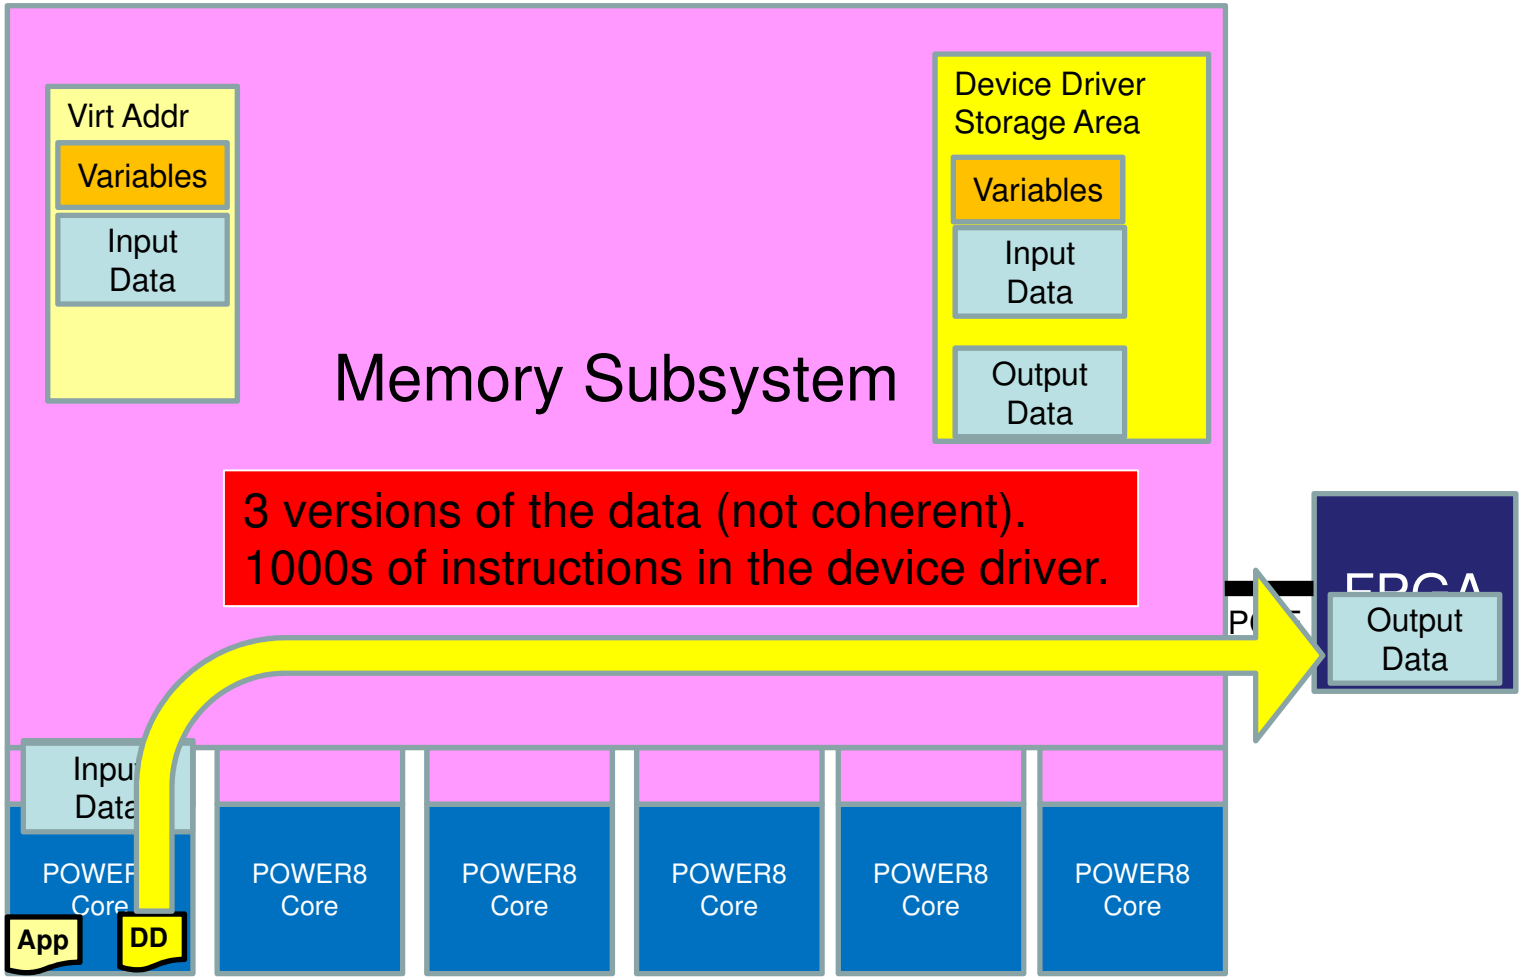

CAPI Coherency

With CAPI, the FPGA shares memory with the cores

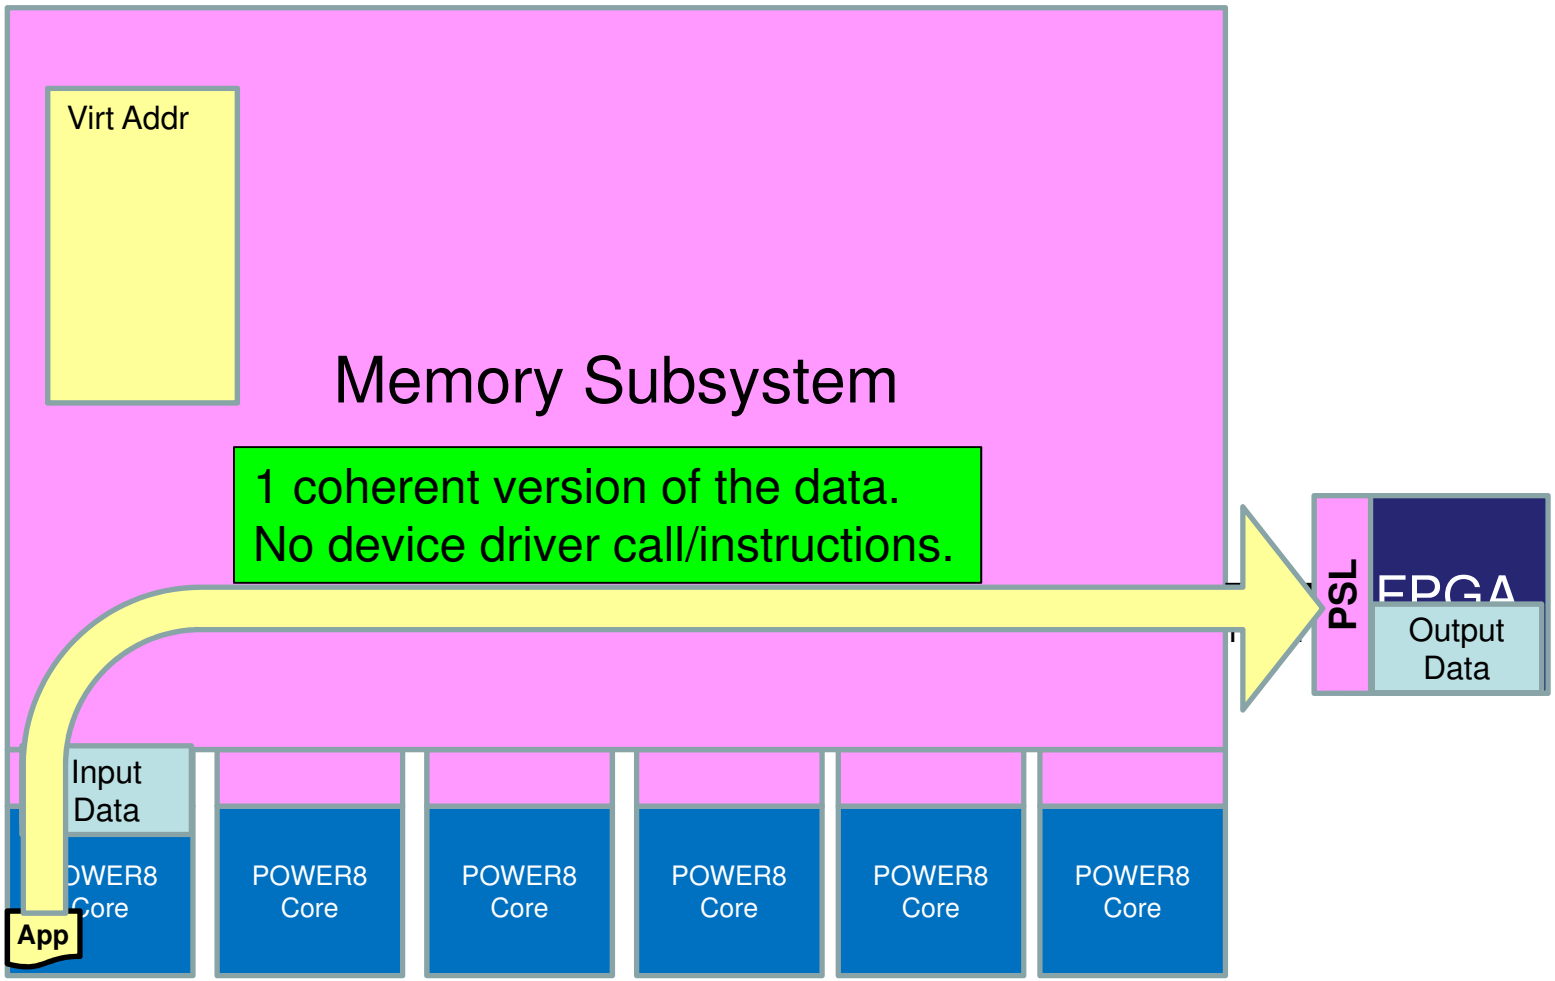

CAPI vs. I/O Device Driver: Data Prep

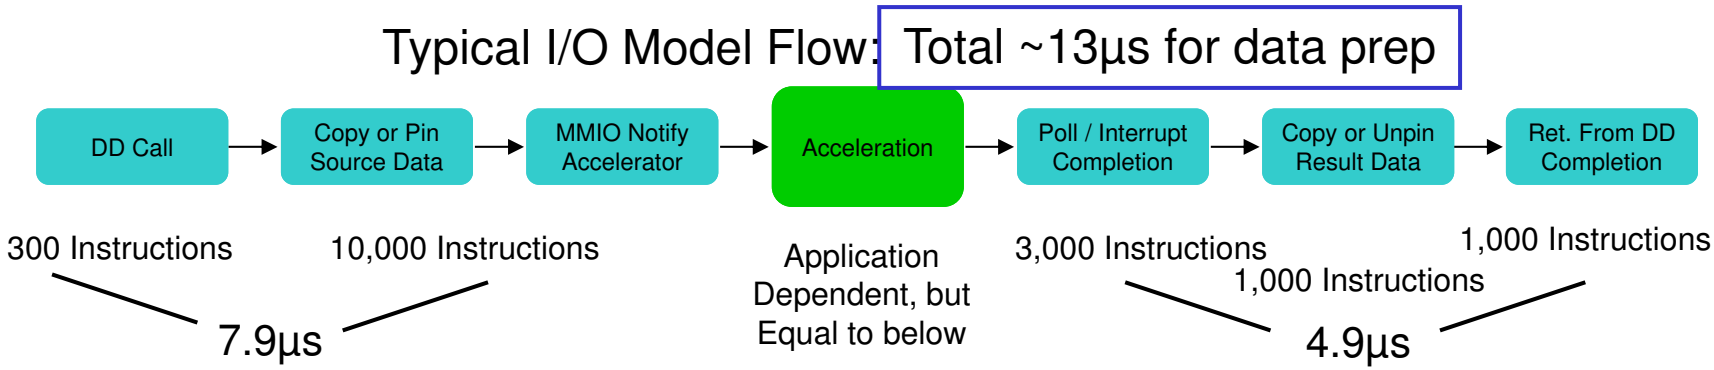

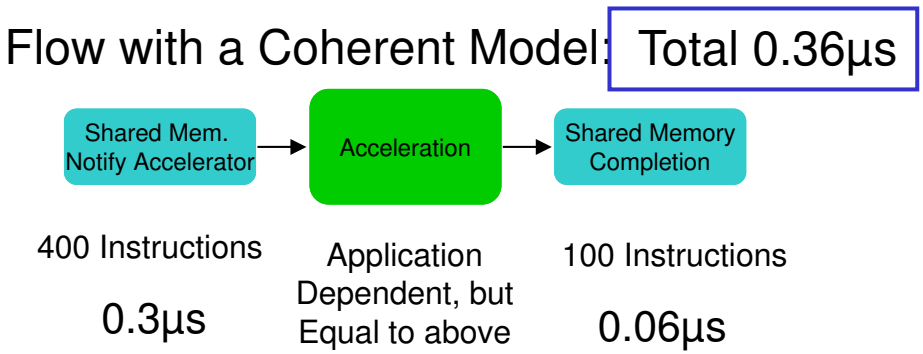

CAPI Differentiation

CAPI vs. I/O or Socket FPGA Solution

IBM Innovation		Customer Impact
FPGA is a peer to the processor -- Caching and translations by PSL	→	Simple Programming paradigm Higher performance
Architecture allows for any kind of FPGA or even an ASIC	→	Flexible solutions Connection to Flash, FC, EN….
Virtualization in the Architecture	→	Applications can share Accelerator

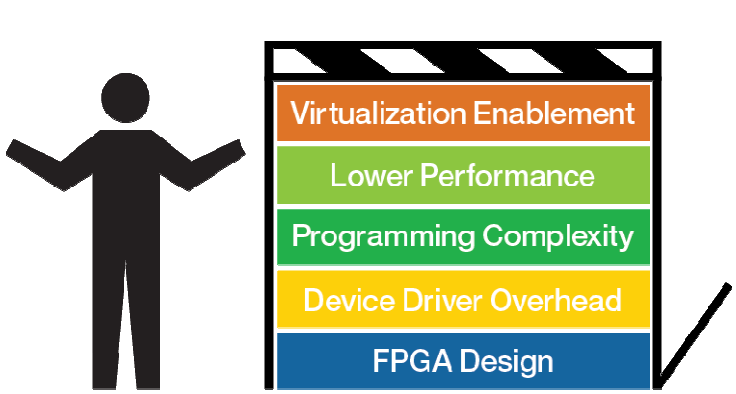

I/O Paradigm

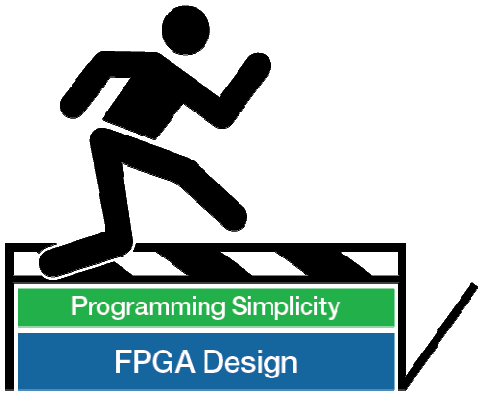

CAPI Paradigm

POWER8 Processor

Let's take a closer look at how IBM Engineers made CAPI work

Technology

- 22 nm SOI, eDRAM, 15 ML 650 mm2

Cores

- 12 cores (SMT8)
- 8 dispatch, 10 issue, 16 execution pipes
- 2x internal data flows/queues
- Enhanced prefetching
- 64 KB data cache, 32 KB instruction cache

Accelerators

- Crypto and memory expansion
- Transactional memory
- VMM assist
- Data move/VM mobility

POWER8 Scale-Out Dual Chip Module

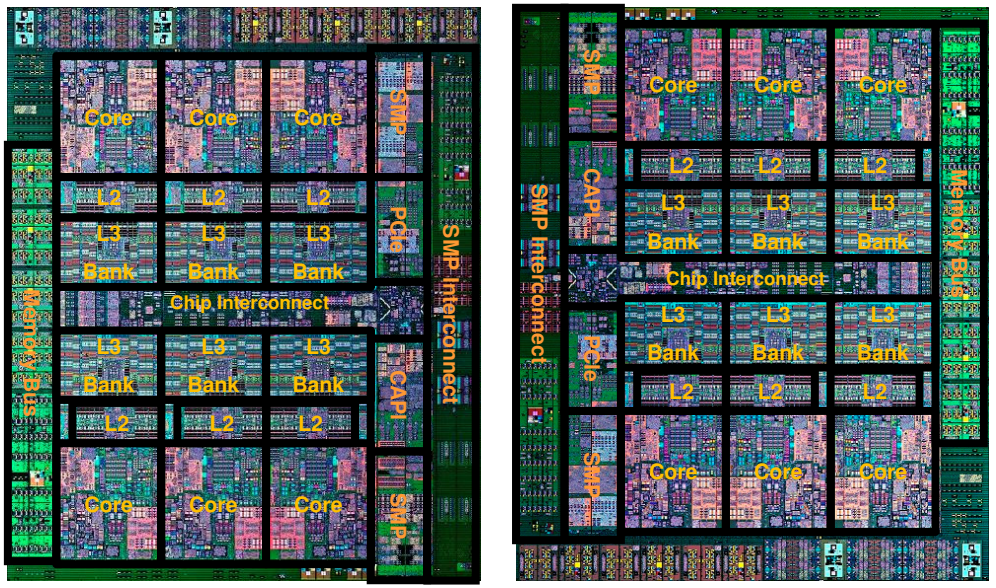

Caches

- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3

Memory

- Up to 230 GB/s sustained bandwidth

Bus Interfaces

- Durable open memory attach interface
- Integrated PCIe Gen3
- SMP interconnect
- CAPI

Energy Management

- On-chip power management microcontroller

How CAPI Works

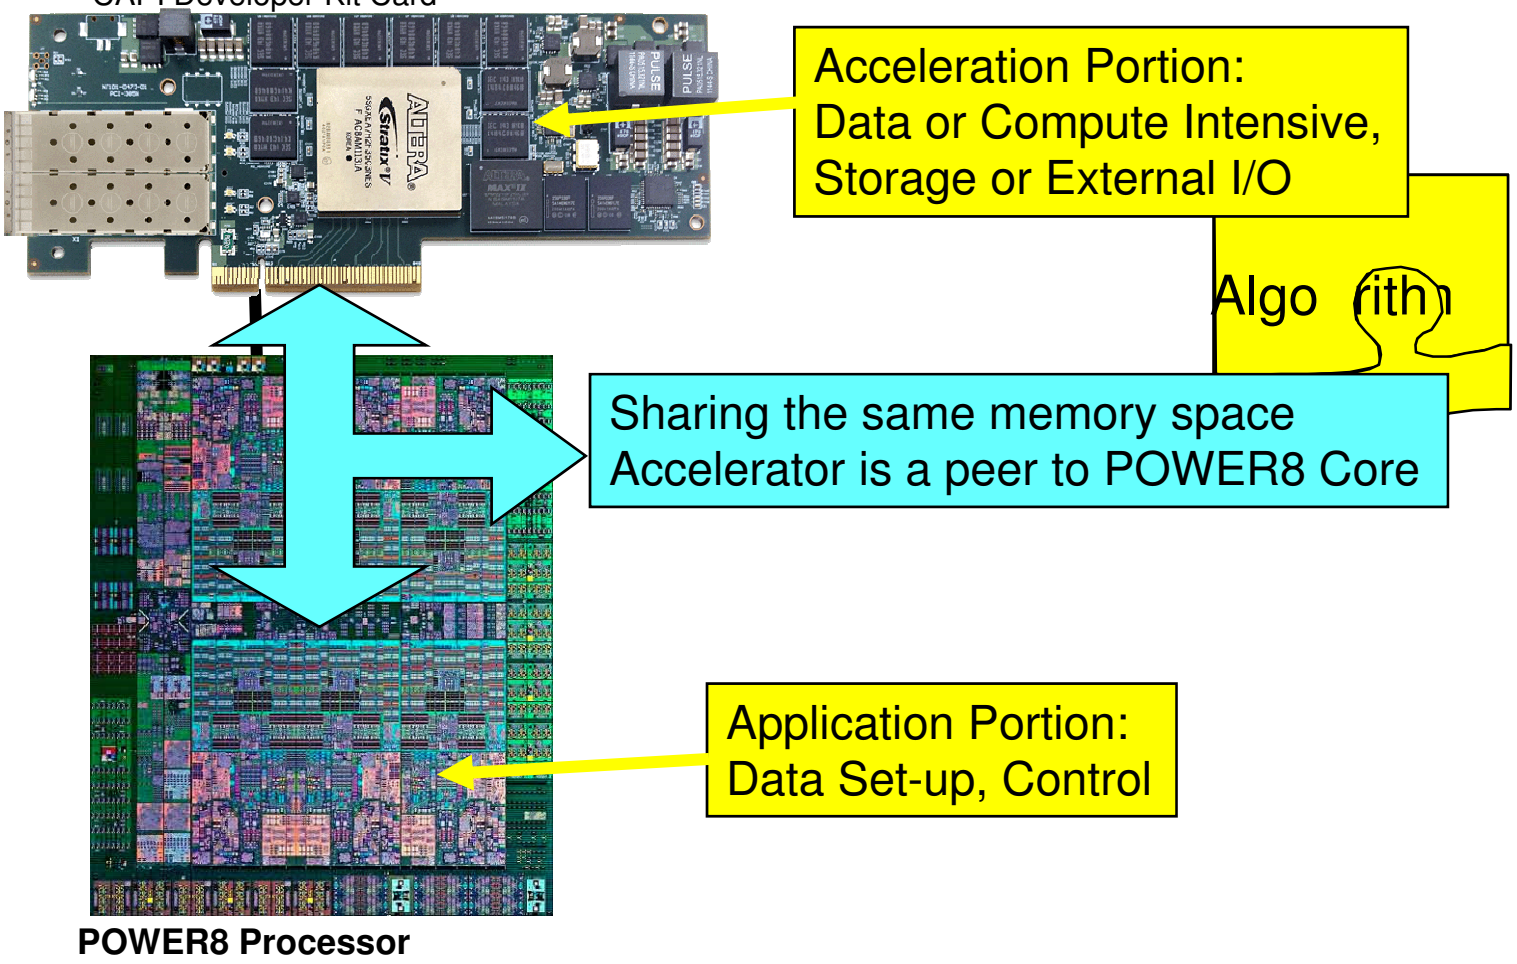

POWER8 with CAPI Cards

Front View

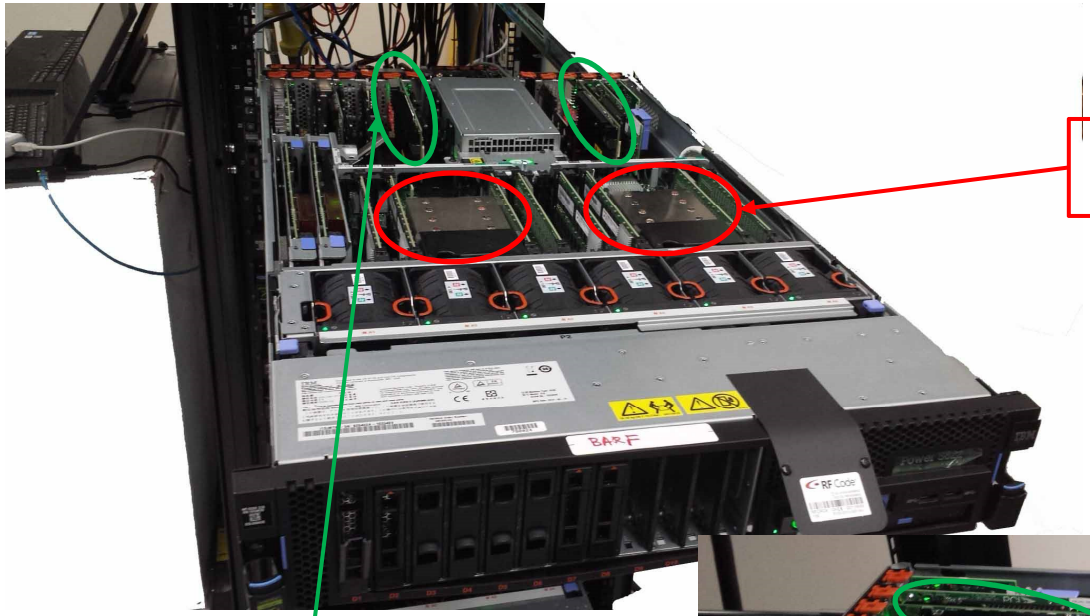

POWER8 Modules

Side View

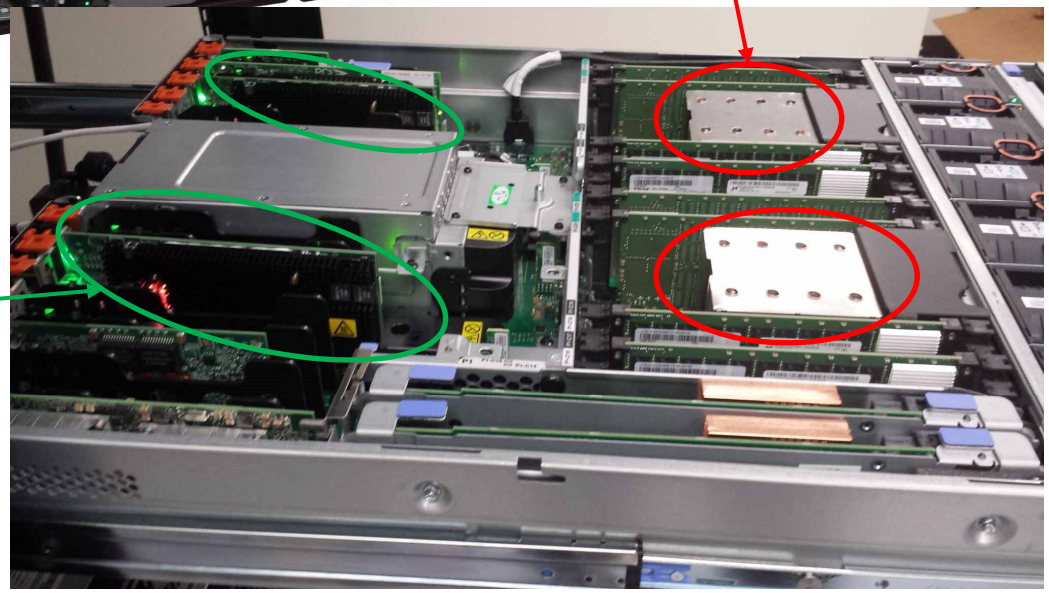

CAPI Dev Kit Cards

Basic concepts of CAPI

CAPI vs. CAPI Solutions

- CAPI is a platform to enable acceleration
- CAPI provides an infrastructure to improve performance of an application through FPGA acceleration
  - Enables customer-defined acceleration within the processor complex
- CAPI allows implementation of a wide range of accelerators to optimally address many different customer challenges
  - Each implementation is a unique CAPI Solution

Platform for Innovation

- A CAPI Solution is a specific implementation of an algorithm that uses an FPGA + application
- A CAPI Solution requires logic designers and programmers to implement the solution
- CAPI Solution Examples:
  - Flash Appliance (IBM Data Engine for NoSQL)
  - MonteCarlo Algorithm

Specific Customer Solution

In Summary

- We compared three paradigms:
  1. Software
  2. IO attached Accelerator
  3. CAPI attached Accelerator

- We showed the components that make CAPI work
  1. POWER8 Hardware: CAPP, PSL, Coherent Memory, PCIE
  2. OS Extensions
  3. Customer's Solution Algorithm and Application

Why Accelerate on CAPI?

- Reasons to consider CAPI Acceleration
  - Higher Performance
    - If your customer has a complex application running on a core, consider CAPI for better performance
    - If your customer already does I/O attached FPGA acceleration, CAPI will simplify their software and provide better performance
  - Lower IT Costs
    - By moving workload to CAPI, your customer will need fewer cores
    - In some cases, such as the IBM Data Engine for NoSQL, CAPI can do the same work with far less infrastructure
  - Lower Power
    - Running acceleration on an FPGA can result in lower power consumption vs. running the application as software on a core

Note:
When considering CAPI for a particular solution, we compare it to:
1. The same solution running as software –OR–
2. The same solution running on an IO attached FPGA

20

CAPI ecosystem partners and consumers

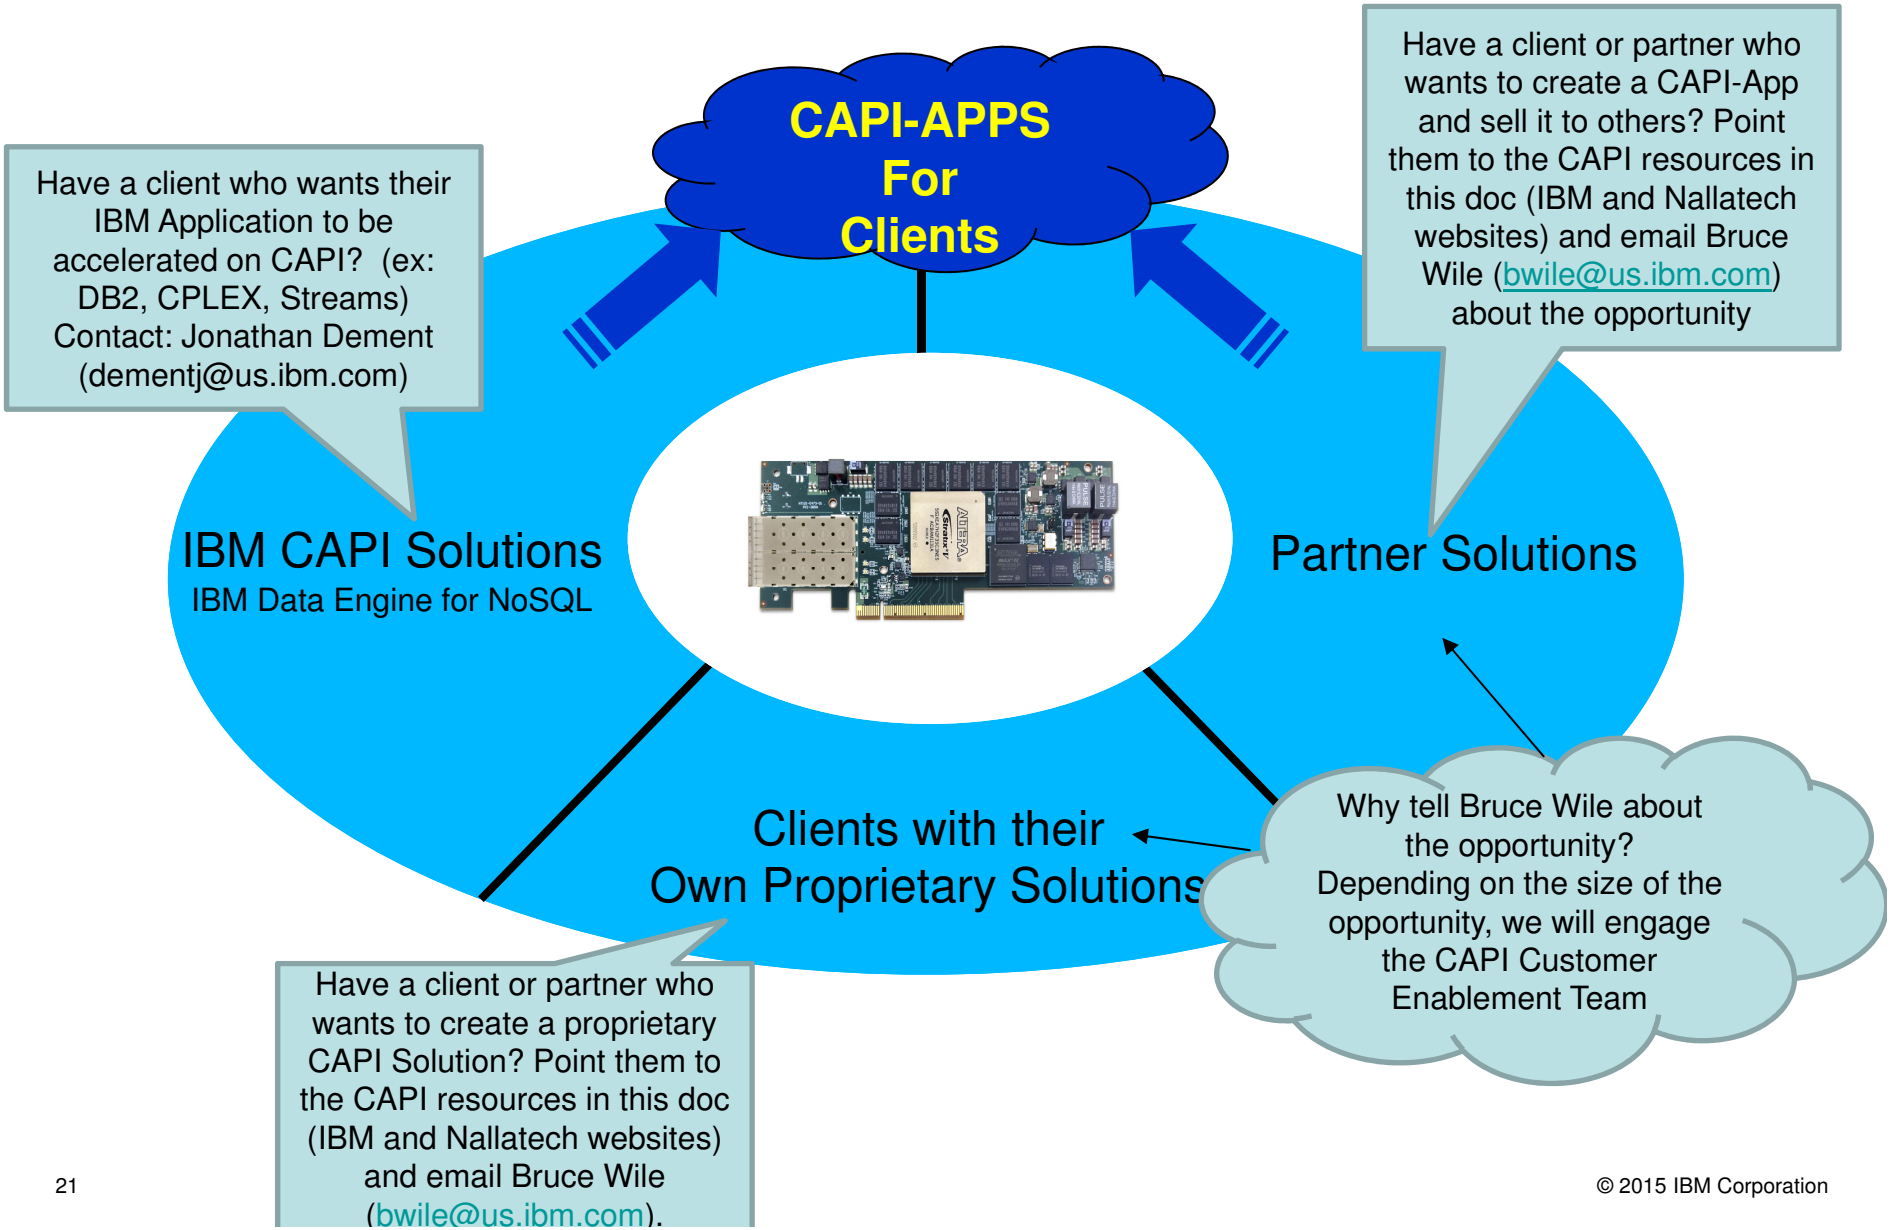

Two Paths into CAPI

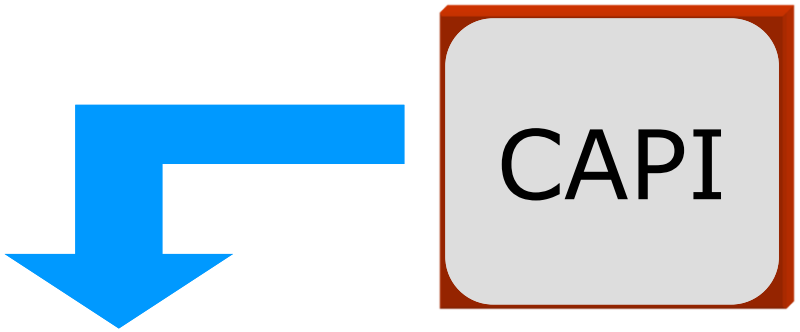

CAPI Developer Kit

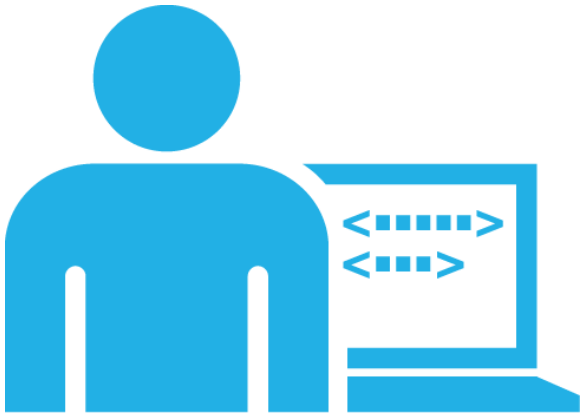

Clients create their own, proprietary business solution.

Nallatech IBM

CAPI Market Solutions

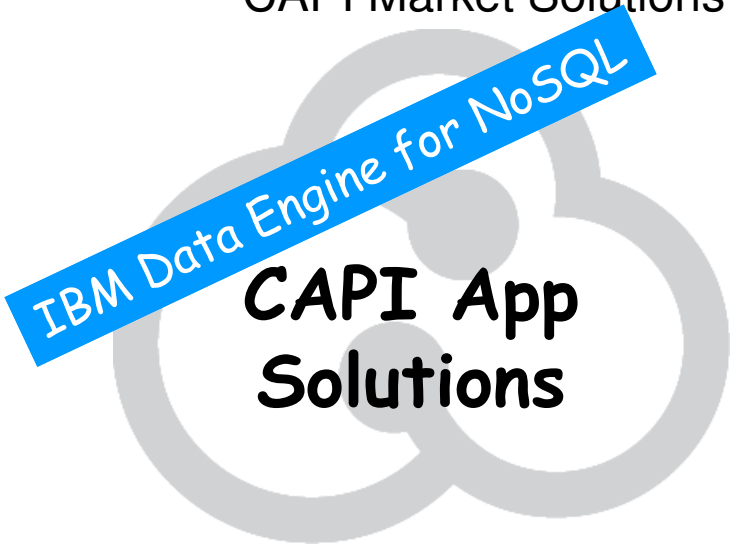

IBM & Partners create business solutions for the CAPI Market. Clients buy pre-packaged solutions from the CAPI Market.

CAPI Solutions

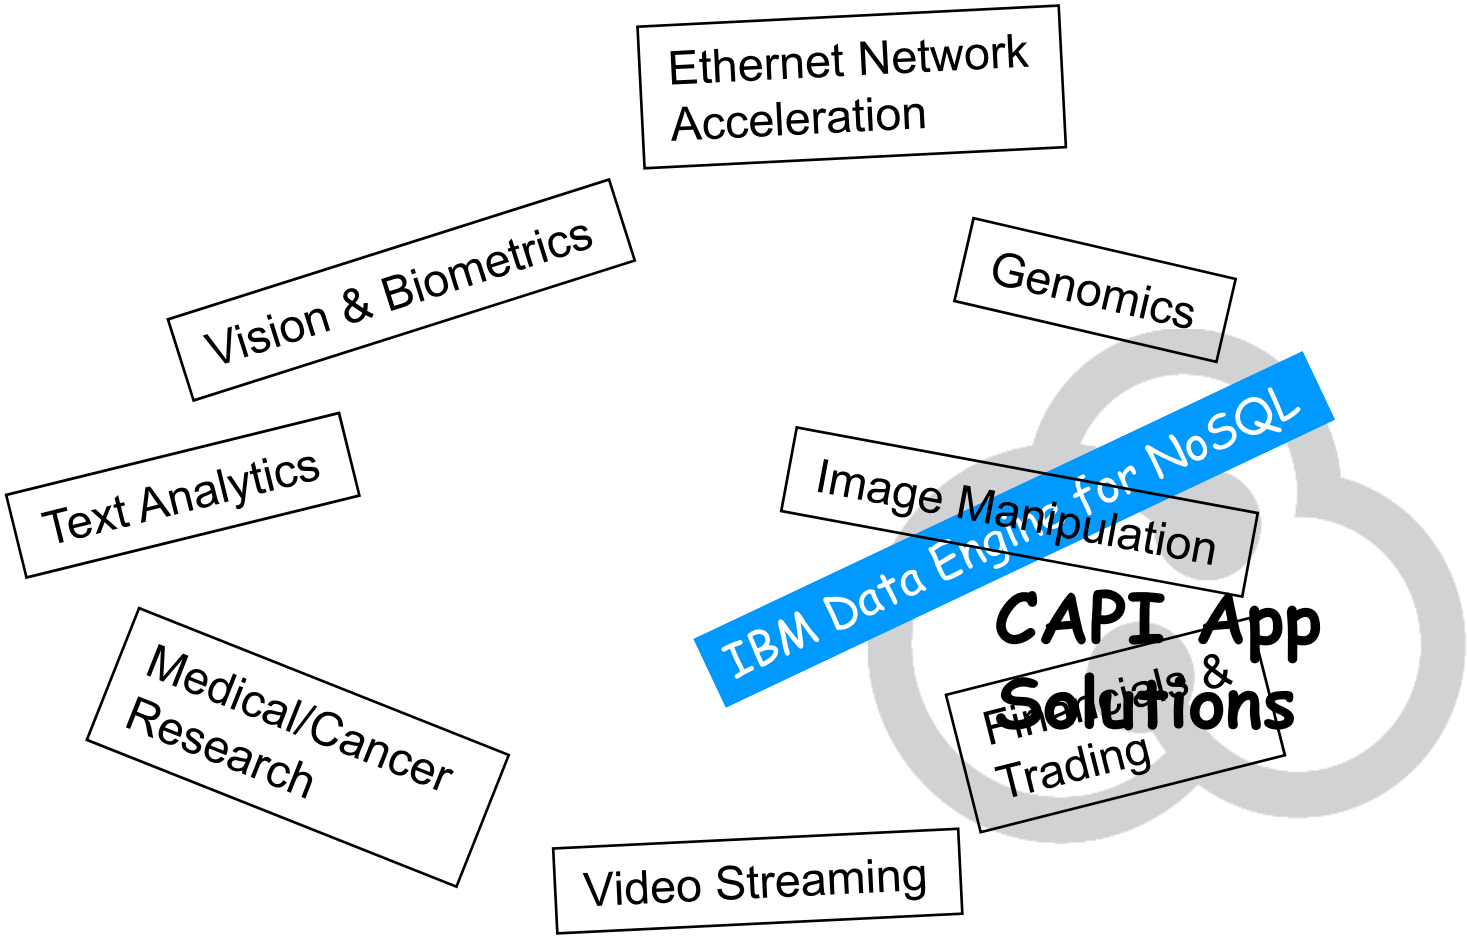

Open Development Driving CAPI Solutions

Power Systems

IBM

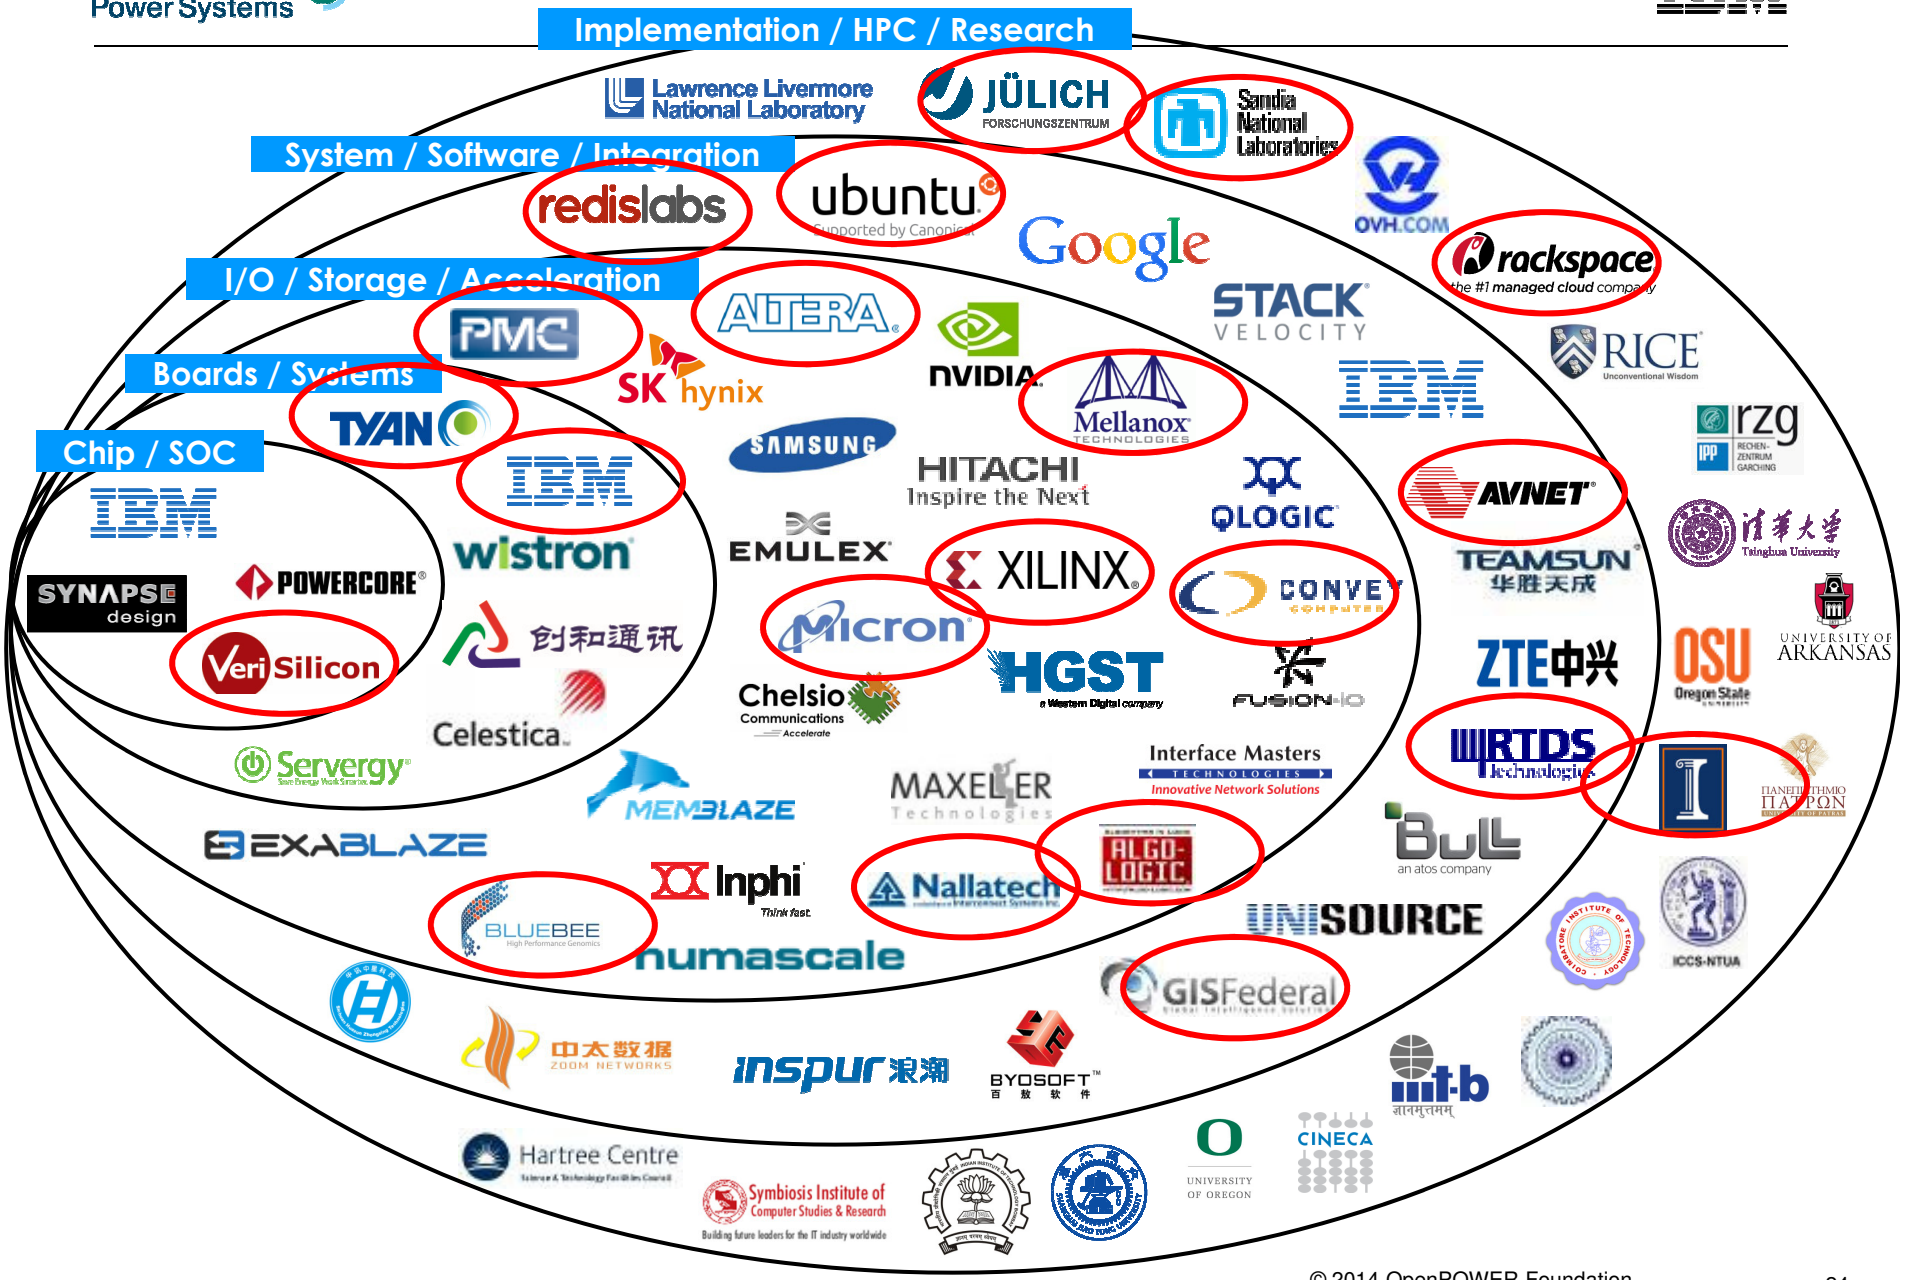

Complete member list at www.openpowerfoundation.org

Potential Markets for CAPI Solutions

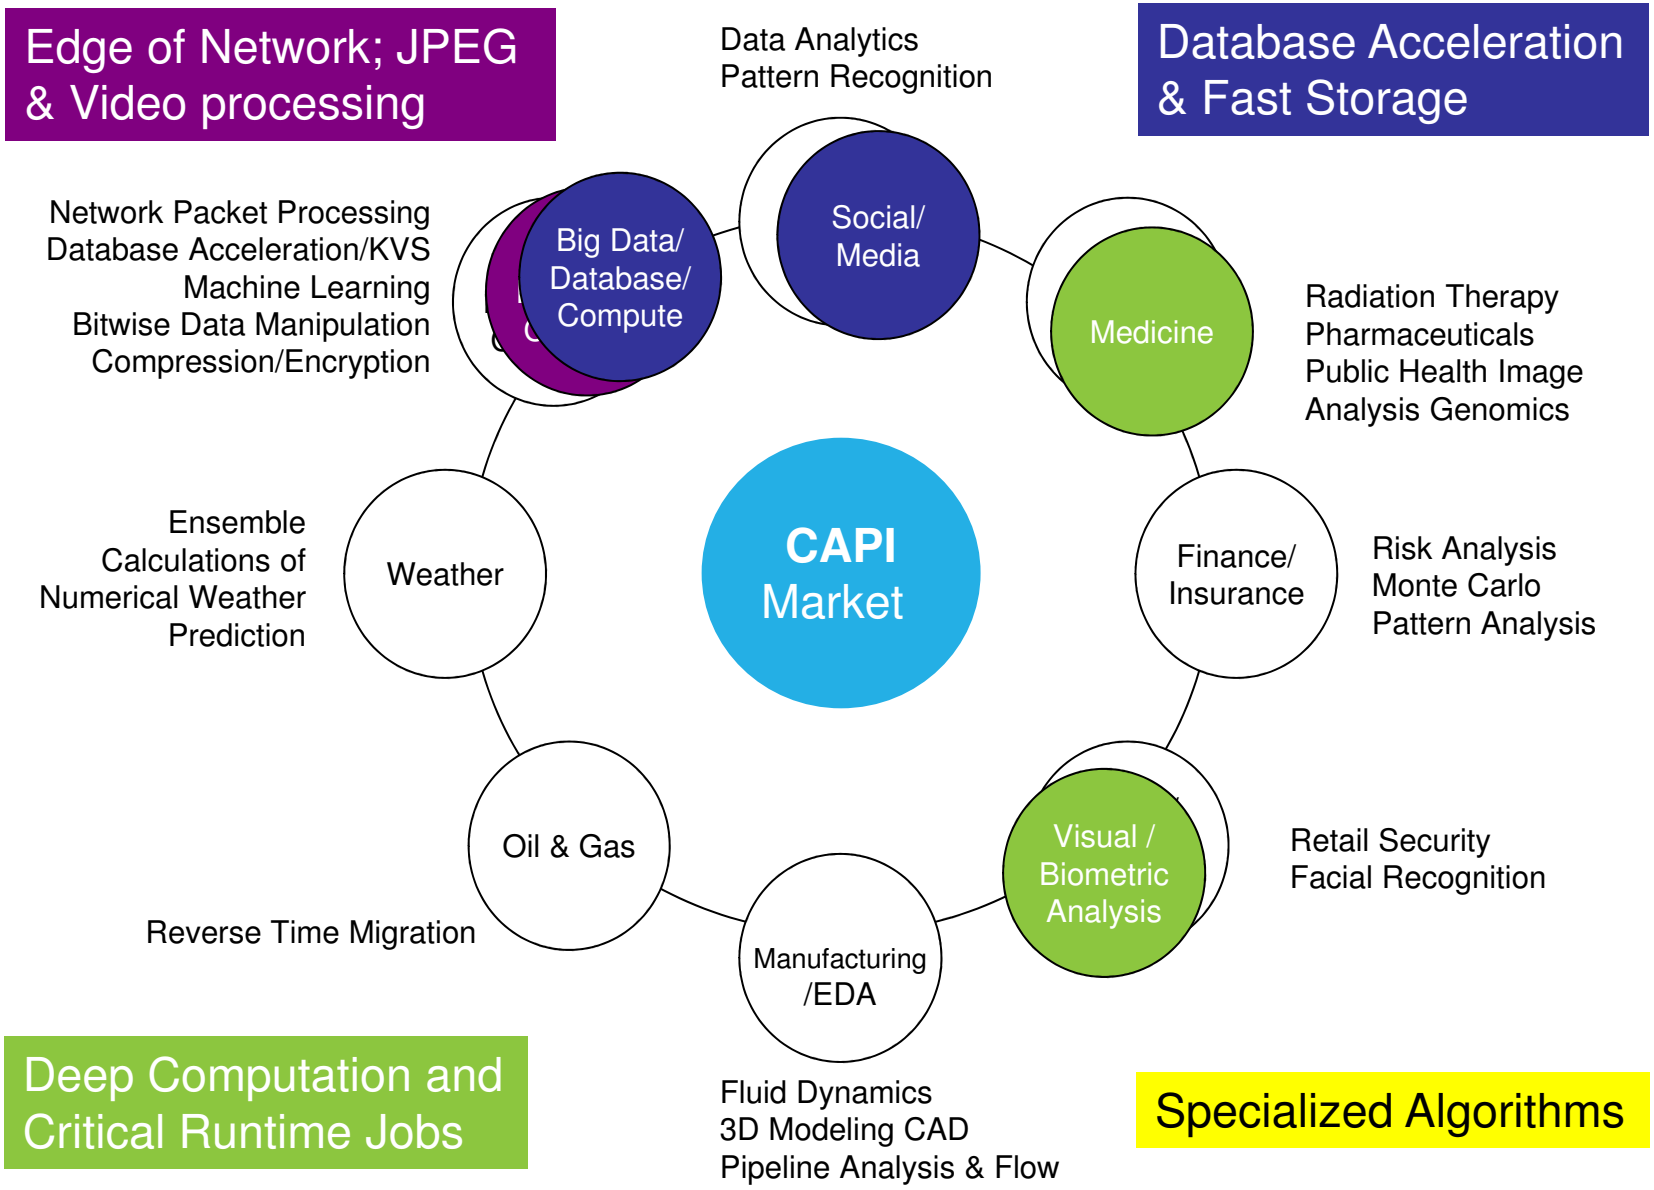

CAPI Availability

- See: http://www.ibm.com/support/customercare/sas/f/capi/home.html
- CAPI Developer Kit
  - Procure through Nallatech

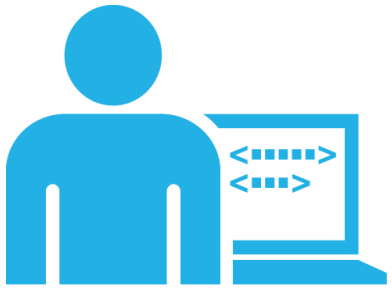

  - For customers considering creating their own CAPI Solution
    - CAPI Decision and Process Guide
  - Requires POWER8 Server
  - Available now
  - See www.nallatech.com/capi

- First CAPI Solution:

IBM Data Engine for NoSQL

  - Procure through IBM
  - GA in early 2015

CAPI Developer Kit

Nallatech's CAPI Developer Kit Contents

	Hardware Components	Software / IP Components	Tools	Documentation
Included in CAPI Developer Kit	Nallatech 385 FPGA Accelerator	IBM CAPI Power Service Layer (PSL) (Encrypted FPGA IP)	Altera Quartus FPGA Tools	White Paper and Decision Guide
	Nallatec JTAG Debug Kit	CAPI Host Support Library (libcxl)	PSL Simulation Engine	CAPI User's Guide
		'Memcopy' Example		385 FPGA Card User Guide
Also Required	Power 8 System (IBM Model 8247-21L or 8247-22L)	CAPI Enabled O/S (initially Ubuntu 14.10 LE from Canonical)	HDL Simulator (i.e. Cadence, Mentor, Synopsis)	

CAPI Developer Kit – FPGA Card

Nallatech's CAPI Developer Kit Contents

	Hardware Components	Software / IP Components	Tools	Documentation
Included in CAPI Developer Kit	Nallatech 385 FPGA Accelerator	IBM CAPI Power Service Layer (PSL) (Encrypted FPGA IP)	Altera Quartus FPGA Tools	White Paper and Decision Guide
	Nallatec JTAG Debug Kit	CAPI Host Support Library (libcxl)	PSL Simulation Engine	CAPI User's Guide
		'Memcopy' Example		385 FPGA Card User Guide
Also Required	Power 8 System (IBM Model 8247-21L or 8247-22L)	CAPI Enabled O/S (initially Ubuntu 14.10 LE from Canonical)	HDL Simulator (i.e. Cadence, Mentor, Synopsis)	

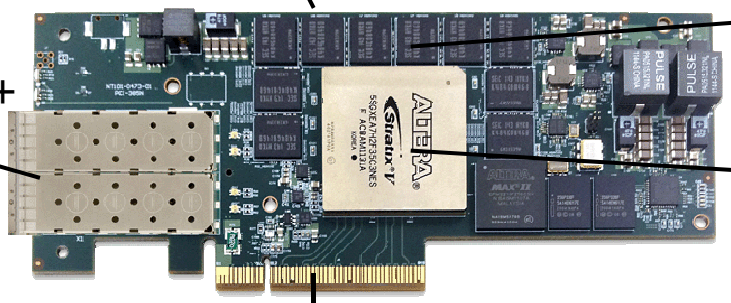

Dual 10G SFP+

2 Banks of SDRAM

Altera Stratix V FPGA

Complete Datasheet

PCI-E Gen3

CAPI Developer Kit

Nallatech's CAPI Developer Kit Contents

	Hardware Components	Software / IP Components	Tools	Documentation
Included in CAPI Developer Kit	Nallatech 385 FPGA Accelerator	IBM CAPI Power Service Layer (PSL) (Encrypted FPGA IP)	Altera Quartus FPGA Tools	White Paper and Decision Guide
	Nallatec JTAG Debug Kit	CAPI Host Support Library (libcxl)	PSL Simulation Engine	CAPI User's Guide
		'Memcopy' Example		385 FPGA Card User Guide
Also Required	Power 8 System (IBM Model 8247-21L or 8247-22L)	CAPI Enabled O/S (Initially Ubuntu 14.10 LE from Canonical)	HDL Simulator (i.e. Cadence, Mentor, Synopsis)	

IBM POWER8™ Server

CAPI Developer Kit

Nallatech's CAPI Developer Kit Contents

	Hardware Components	Software / IP Components	Tools	Documentation
Included in CAPI Developer Kit	Nallatech 385 FPGA Accelerator	IBM CAPI Power Service Layer (PSL) (Encrypted FPGA IP)	Altera Quartus FPGA Tools	White Paper and Decision Guide
	Nallatec JTAG Debug Kit	CAPI Host Support Library (libcxl)	PSL Simulation Engine	CAPI User's Guide
		'Memcopy' Example		385 FPGA Card User Guide
Also Required	Power 8 System (IBM Model 8247-21L or 8247-22L)	CAPI Enabled O/S (initially Ubuntu 14.10 LE from Canonical)	HDL Simulator (i.e. Cadence, Mentor, Synopsis)	

CAPI Developer Kit

Nallatech's CAPI Developer Kit Contents

	Hardware Components	Software / IP Components	Tools	Documentation
Included in CAPI Developer Kit	Nallatech 385 FPGA Accelerator	IBM CAPI Power Service Layer (PSL) (Encrypted FPGA IP)	Altera Quartus FPGA Tools	White Paper and Decision Guide
	Nallatec JTAG Debug Kit	CAPI Host Support Library (libcxl)	PSL Simulation Engine	CAPI User's Guide
		'Memcopy' Example		385 FPGA Card User Guide
Also Required	Power 8 System (IBM Model 8247-21L or 8247-22L)	CAPI Enabled O/S (initially Ubuntu 14.10 LE from Canonical)	HDL Simulator (i.e. Cadence, Mentor, Synopsis)	

http://www.ibm.com/support/customercare/sas/f/capi/home.html

Capturing the growth in Big Data

- Growth of NoSQL solutions is explosive
  - In Memory requirements drive large infrastructures and cost
  - Deployment complexity limits growth
- IBM CAPI-Flash delivers a new size/performance price point
  - Equivalent end to end performance
  - Significantly lower deployment and operational costs
  - Reduce infrastructure complexity
- 100% Redis Compliant
  - Your redis applications just work!
  - No lock in
  - You choose the amount of real memory vs flash memory and we handle the rest
  - Huge savings 3x cheaper and nearly the same performance
    (for well behaved access patterns)

IBM Data Engine for NoSQL

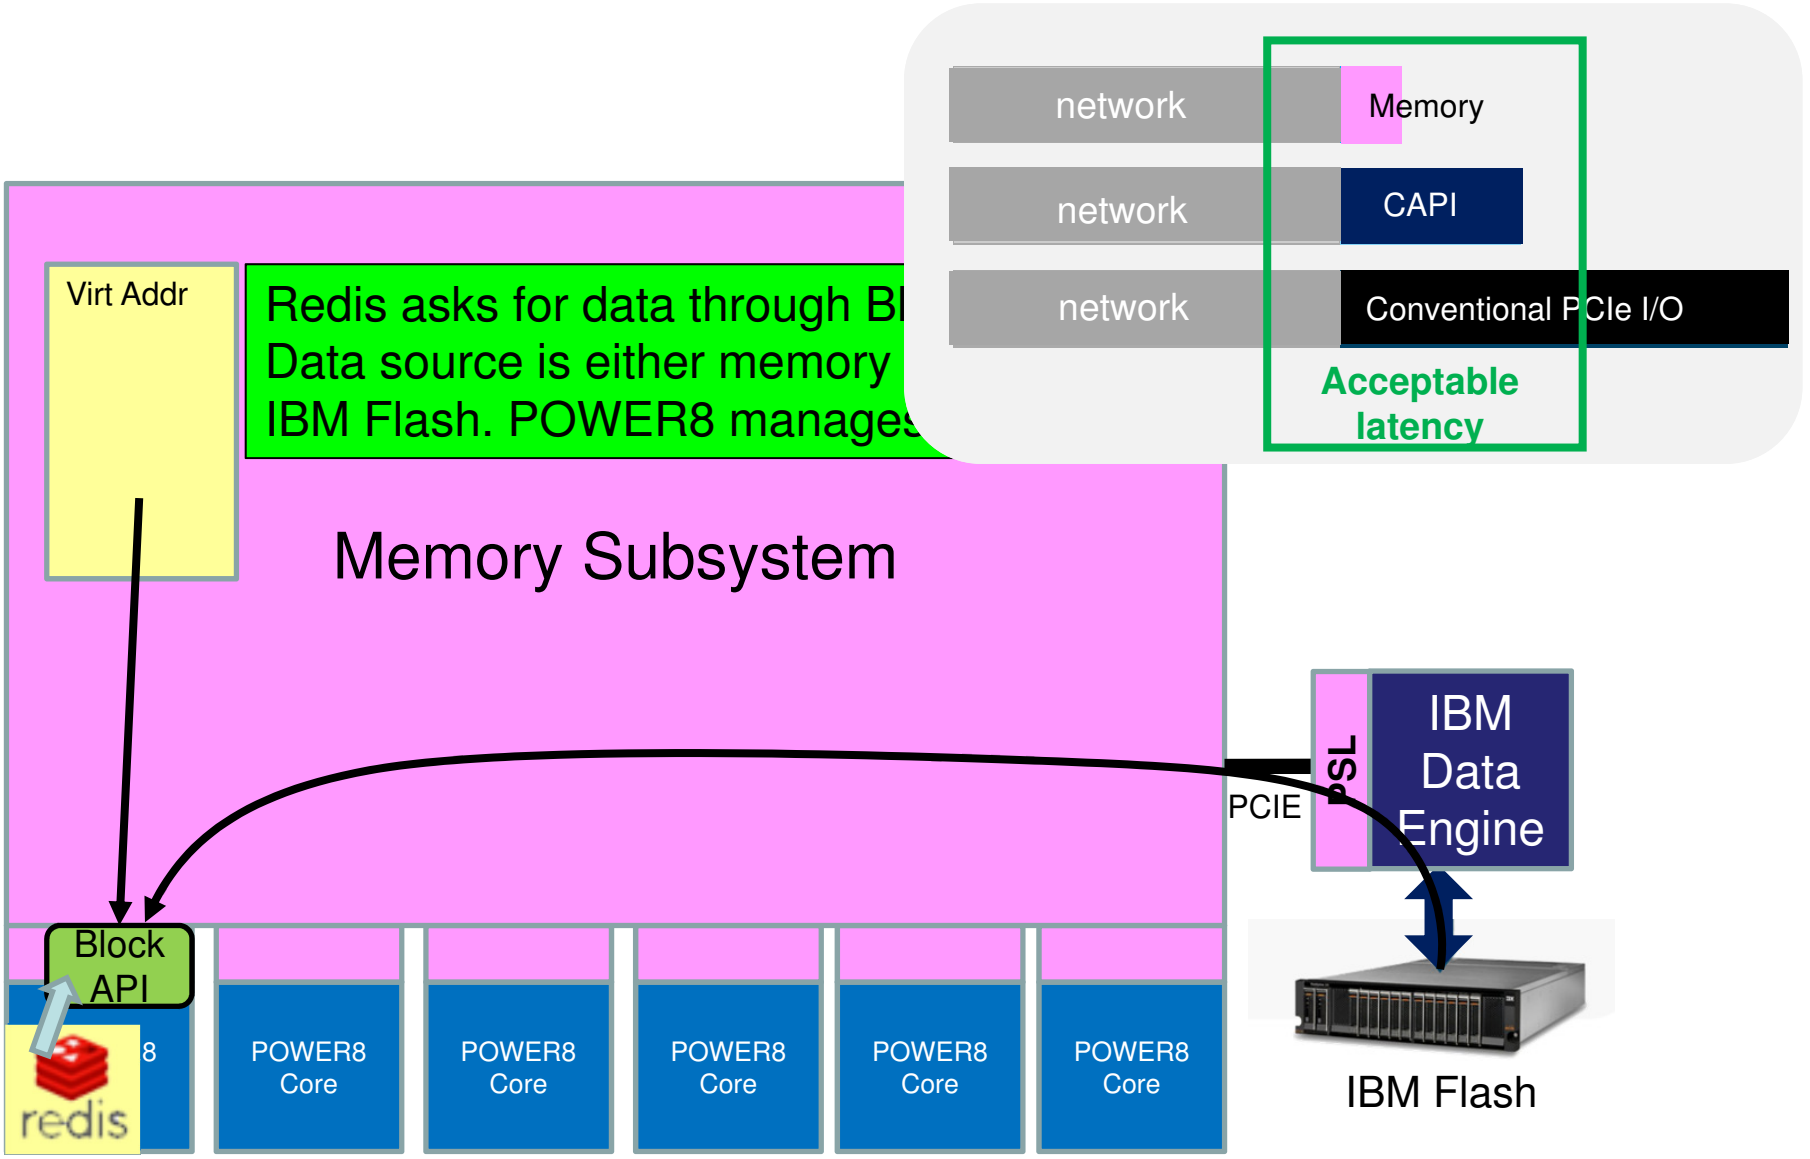

Innovative “In-Memory” NoSQL/KVS

IBM Data Engine for NoSQL

Today’s NoSQL in memory (x86)

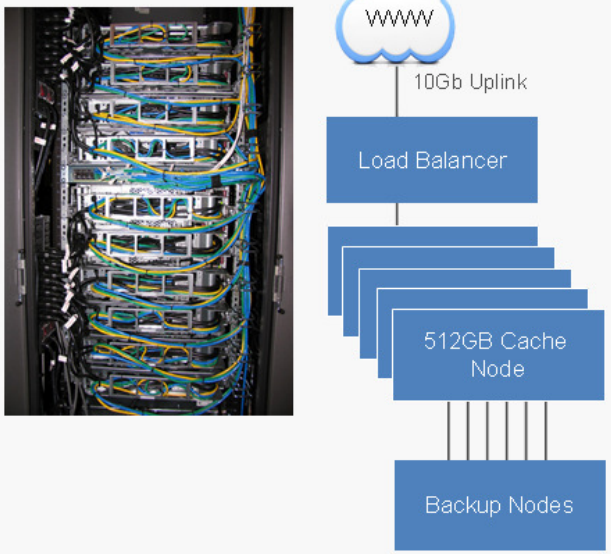

4Q14 - Differentiated NoSQL (POWER8 + CAPI Flash solution)

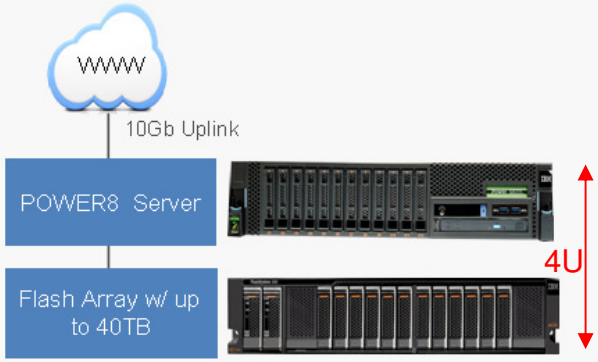

24:1
Reduction in infrastructure

2.4x
Price reduction

12x
Less Energy

6x
Less rack space

40TB of extended memory

Target LoB / solution architects and MSPs

- Supporting or building mobile/web/social apps
- Leveraging Key Value Store (KVS) for fast lookups
- Require high performance in-memory data access

Power + CAPI Flash Advantage

- 24:1 physical server consolidation
- 6x less rack space (2U server+2U Flash vs. 24 1U servers)
- Regain infrastructure control
- Dramatically reduce costs to deliver services

Demonstrating the Value of CAPI Attachment

IBM Data Engine for NoSQL

Identical hardware with 2 different paths to data

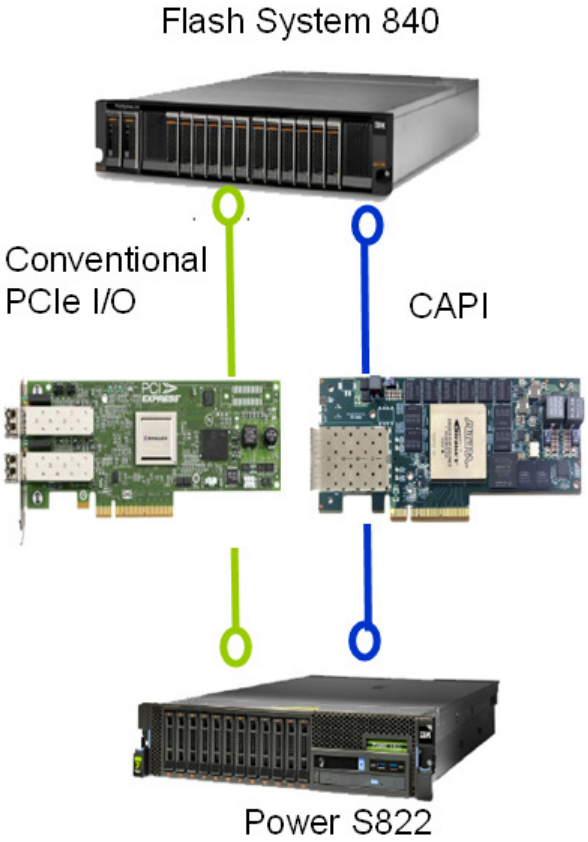

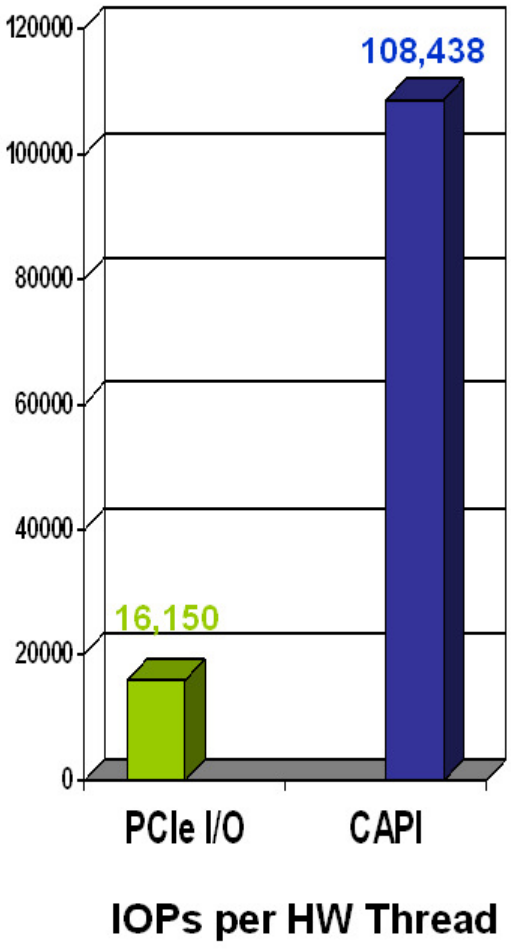

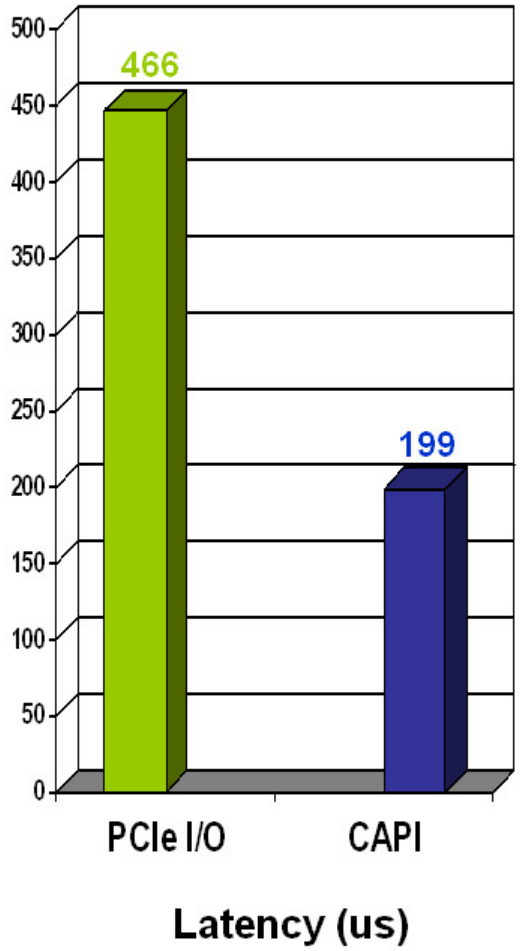

Ways to get Started:

- Move your REDIS or MEMCACHED application over transparently
  - BigRedis and Memcached from Redis Labs is 100% compliant with current client APIs.
  - Your redis applications just work!
- Need a different NoSQL?
  - Talk to IBM as we are working with other NoSQL providers and you may be able to join a Proof of Concept when available
- Want to move your application/service directly to or KVS or Block APIs?
  - NDA release of APIs
  - Development Tools & Education available
    - Dev Kit available for Power8 Systems
    - Use time on Systems available online
    - Engage IBM for a Proof of Concept project

Workloads to Innovate

- Start with what FPGAs are good at: Embarrassingly Parallel Problems
- Combine with CAPI strengths:
  - Ease of programming
  - Lack of device driver
  - Shared memory & caching (host to accelerator communication)
- What do you get:
  - Bitwise data manipulation (e.g. Deep Compression)
  - Pattern recognition
  - Encryption
  - Monte Carlo
    - Statistical modeling for complex predictions
  - Image Analytics & Biometrics
    - Facial recognition
    - Feature detection (e.g. cancer)
  - Network Packet Processing & Inspection
  - Bioinformatics (e.g. Sequence alignment)
  - Reverse time migration (Oil & Gas)
  - Ensemble Calculations of Numerical Weather Prediction
  - Machine Learning
  - And on and on

IBM Accelerated GZIP Compression

What it is:

- An FPGA-based low-latency GZIP Compressor & Decompressor with **single-thread** througput of ~2GB/s and a compression rate significantly better than low-CPU overhead compressors like snappy.

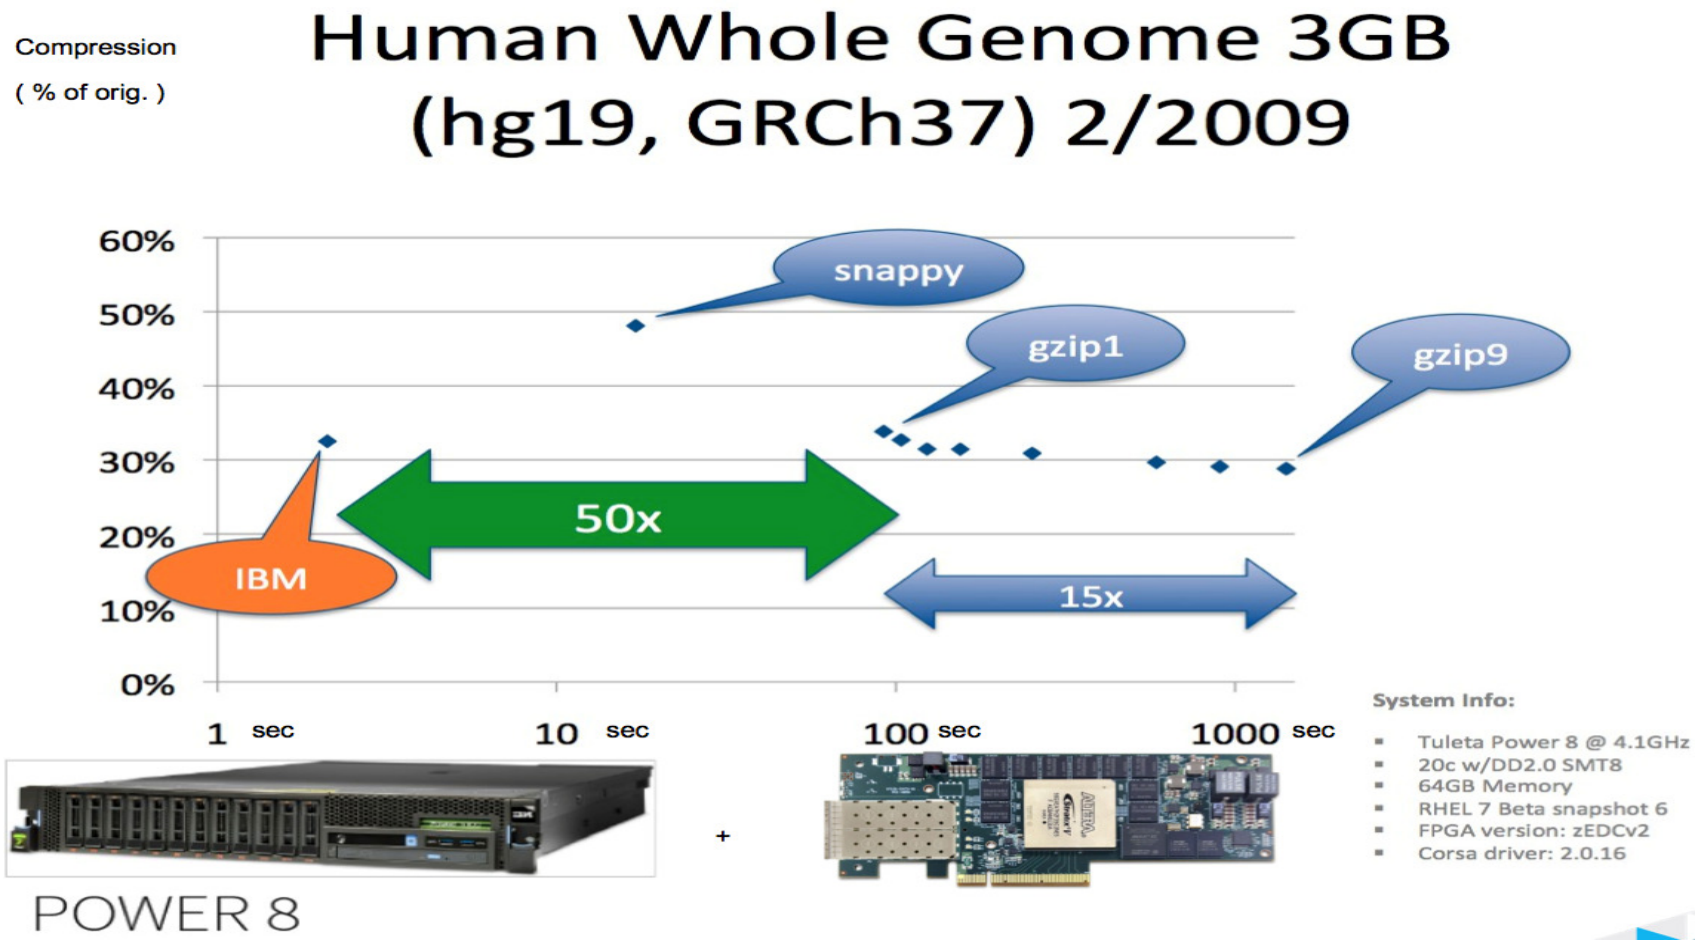

IBM Accelerated Text Processing

What it is:

- A compiler/runtime system for accelerating text analytics on a shared-memory CPU-FPGA

Results

- Big Speedup vs. Multithread SW

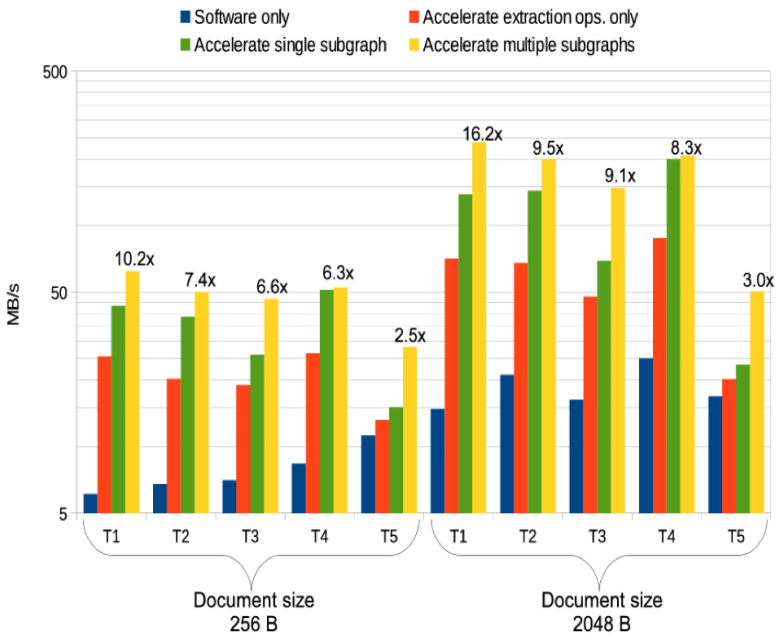

To appear @:

- Hot Chips 2014

AQL
- rule language
- SQL-like synta

systemT optimizer

Compiled operator graph

systemT runtime
Java + FPGA

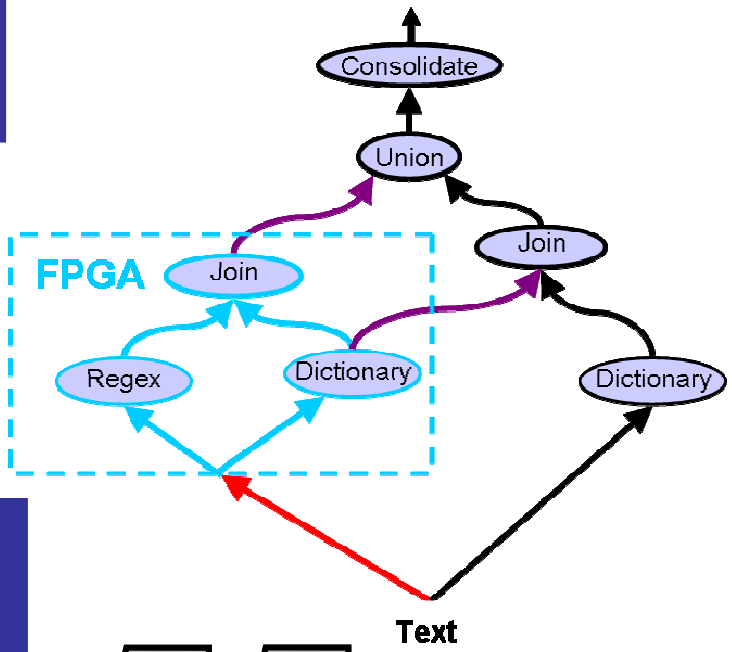

Annotations

For years, Microsoft Corporation CEO Bill Gates was against open source. But today he appears to have changed his mind. "We can be open source"

FPGA Image & Video Processing

Power Systems

IBM

Information Extraction → **Object Recognition**

Edge Detection, Feature Extraction, Segmentation

Template Matching

Goal

Extract ***relevant information*** from input image to enable ***object recognition***

Information located where pixels change color (edges, blobs)

- *Intrinsic properties of objects*
- *Object boundaries*

Approach

Design fully-pipelined FPGA architectures ➔ streaming application

Real-time, low-power, onboard image processing solution

- *Sobel and Canny: extract contours/edges*
- *SURF: extract scale & rotation-invariant features*

OpenCV

Motivations

Applications requiring edge detection & feature extraction span a wide range of domains

- *Computer/Machine Vision: Tracking, Object Recognition & Navigation*
- *General image proc.: Compression*
- *Quality Control: Unsupervised Defect Identification*
- *Medical Imaging: Analysis + Diagnosis & Computer Guided Surgery*

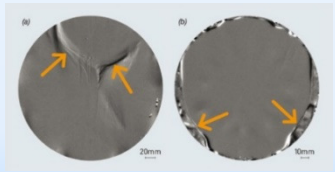

Directions for CAPI Ideas

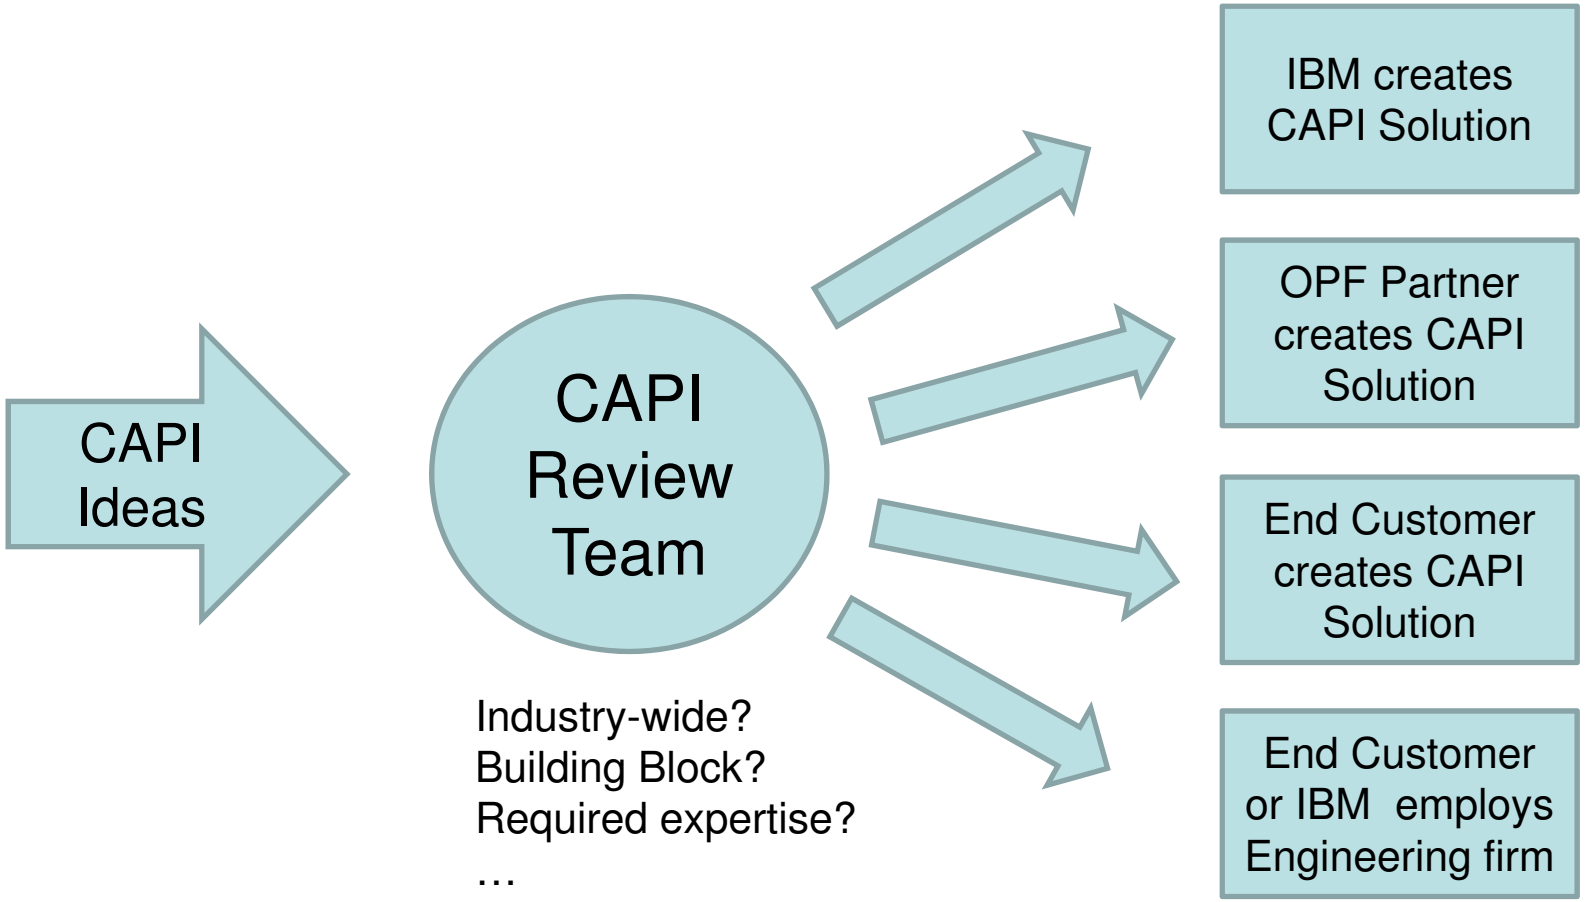

Printed in the United States of America September 2014

IBM Systems and Technology Group
2070 Route 52, Bldg. 330
Hopewell Junction, NY 12533-6351

The IBM home page can be found at **ibm.com®.**

Version 1.0
29 September 2014—IBM Confidential